Payot, A. D., Kedward, L. J., Rendall, T., & Allen, C. B. (2019). Optimisation of Multi-Modal Aerodynamic Shape and Topology Problems. In *AIAA Scitech 2019 Forum: 7-11 January 2019, San Diego, California* (AIAA Scitech 2019 Forum). American Institute of Aeronautics and Astronautics Inc. (AIAA).
<https://doi.org/10.2514/6.2019-1206>

Peer reviewed version

License (if available):
Other

Link to published version (if available):
[10.2514/6.2019-1206](https://doi.org/10.2514/6.2019-1206)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the accepted author manuscript (AAM). The final published version (version of record) is available online via AIAA at <https://doi.org/10.2514/6.2019-1206> . Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Optimisation of Multi-Modal Aerodynamic Shape and Topology Problems

A. D. J. Payot ^{*}; L. J. Kedward [†]; T. C. S. Rendall [‡]; C. B. Allen [§]

Department of Aerospace Engineering, University of Bristol, Bristol, UK

This paper presents further results on the development of a combined multi-fidelity shape and topology optimisation framework for aerodynamic optimisation. The combined framework comprises: a multilevel subdivision shape parameterisation used in combination with an adjoint solver and gradient-based optimiser for robust high-fidelity local shape optimisation; and a restricted snake volume of solid parameterisation and differential evolution search algorithm for flexible coverage of a global topological design space. Previous work has demonstrated the efficacy of extending a topological design space by including an efficient local shape optimisation, and the challenges which need to be addressed to do so. Importantly, the work has incorporated the two approaches such that the geometric flexibility of the restricted snake volume of solid method in representing arbitrary topology, has been combined with the efficient and high fidelity precision provided by local shape optimisation. As a result significant performance enhancements have been brought to the topology optimisation problem. The combined framework is benchmarked on a challenging constrained supersonic drag minimisation problem exhibiting multi-body solutions, discontinuities and multi-modality, and significant improvements are demonstrated compared to the individual local and global methods applied separately. In this paper, test cases with additional topological complexity are tackled with focus given to the multi-modality of the aerodynamic objective in the combined shape and topology design space.

I. Introduction and Background

Increases in computational power and improvements in computational fluid dynamics (CFD) tools have created the possibility of using CFD-based optimisation in industrial design. By allowing a systematic and unbiased exploration of a design space, optimisation methods can be used to expand a designer's understanding of the problem being tackled, allowing better overall aerodynamic performance. As designers look to improve performance, aircraft manufacturers are turning increasingly to numerical optimisation. Frameworks for aerodynamic optimisation require the integration of parameterisation methods, mesh generators and flow solvers with optimisation methods. The tendency in this has been to use a modular approach by integrating established modelling and CFD packages with existing optimisers.

Current aerodynamic shape optimisation (ASO) methods must choose between efficient and robust recovery of a local minima, and design space exploration and the search for global minima. The answer to this dilemma is guided by the ASO community's current understanding of the multi-modality of the problem it faces. Evidence suggests that most well posed aerodynamic shape optimisation problems are, in the worst cases, weakly multi-modal. This has meant that the focus has been on developing parameterisations and optimisers which can rapidly converge on a local minimum. What is true for shape optimisation is not applicable to aerodynamic topology optimisation. Allowing the number of external bodies in the flow to change is likely to introduce a number of local minima which can no longer be tackled by current methods. Work within the research group of the authors has so far focused on developing some of the most effective parameterisation methods for two-dimensional aerofoils [1] and new parameterisations which allow topology optimisation for aerodynamics [2–4]. This work aims to integrate these two approaches to create an optimisation framework capable of efficiently exploring the multi-modal design spaces of topology optimisation cases with intricate constraints.

The complexity of parameterisation methods arises from the different origins of optimisation methods and CFD processes. Optimisation methods are mathematical algorithms devised to find the extrema of functions, and have rigorous mathematical underpinnings. Meanwhile, CFD originated from the need to evaluate the aerodynamic properties of potential designs. The translation of the mathematical formulations used by optimisers into the geometric designs used by CFD is a complex problem with implications on the efficiency and effectiveness of optimisation frameworks. To

Copyright © 2019 by Laurence Kedward; Alexandre Payot

^{*}PhD Student, AIAA Student Member, a.payot@bristol.ac.uk, Bristol, BS8 1TR, UK

[†]PhD Student, AIAA Student Member, laurence.kedward@bristol.ac.uk, Bristol, BS8 1TR, UK

[‡]Lecturer, AIAA Member, thomas.rendall@bristol.ac.uk, Bristol, BS8 1TR, UK

[§]Professor of Computational Aerodynamics, AIAA Senior Member, c.b.allen@bristol.ac.uk, Bristol, BS8 1TR, UK

be effective parameterisation methods need to be both compact (need few design variables) and should not artificially limit the geometric space that can be represented [5].

Earlier developments in the field of parameterisation for aerodynamics have yielded a wealth of different methods for the representation of aerodynamic designs. Parameterisation methods can be separated broadly in two categories: constructive and deformative methods. Constructive methods completely define the geometry from the set of design variables; these include B-Spline and polynomial interpolation [6] in general, and CST [7] and PARSEC [8] in particular. Deformative methods by comparison define a set of modifications to a baseline geometry; notable among these are the Hicks-Henne bump functions [9], Singular Value Decomposition (SVD) deformation modes [10, 11] and Free-Form Deformation (FFD) methods [12, 13]. While most parameterisations presented here can be extended to three dimensions, their capability varies widely. In three dimensions the most common is to use FFD deformation methods as these can be adapted to work directly on an existing mesh. Previous systematic investigations by Vassberg et al. [5, 14] have highlighted the impact of dimensionality on the drag minimisation of a standard test case, showing the importance of geometric flexibility while maintaining a compact set of design variables. Work by Castonguay and Nadarajah [15], and more recently by Masters et al. [16, 17] has compared the impact of established parameterisation methods on geometric flexibility, pressure distribution recovery and optimal drag results. These studies show that effective parameterisation methods will require few design variables while providing smooth control of the aerodynamic profile. Smooth control is achieved when a small change in the numerical representation must lead to a similarly small change in the represented geometry. This requirement, resulting from the expense of optimisers to converge in large design spaces, is traded off with the need to not artificially restrict the scope of geometries that can be represented [5]. Most parameterisation methods to date have focused on producing smooth designs with small numbers of design variables. One key geometric restriction that affects all parameterisation methods is the inability to transition between topologies. What this means is that no conventional aerodynamic optimisation framework is currently capable of exploring the number of aerodynamic bodies with a single set of design variables.

In structural design the benefits of exploring different topologies is key to generating light-weight and efficient structures. The field of numerical topological structural optimisation has been an active field of research for the last 30 years and it has recently seen industrial application on the Boeing CH-47 Chinook and the Airbus A380; it allowed a weight reduction of 17% of underfloor beams compared to a conventional structural optimisation method [18] on the CH-47 and weight reduction of the leading edge droop ribs on the A380. This effort in the Finite Elements Analysis (FEA) community has led to parameterisation methods able to represent complex topologies with a single set of design variables [19–21]. The use of volume information to represent geometries has seen widespread use in the structural topological optimisation community.

The justification for topological optimisation is straightforward in structural applications, from truss space-frames to honeycomb designs, there are a wide range of possible engineering structures; furthermore a structural member's impact is readily summarised to a set of interactions at its boundary. The formulation of current FEA code is Lagrangian which allows changing topologies within the modelling architecture without re-meshing the object under analysis. The possibility to reduce designs to a set of external interactions and the Lagrangian formulation of CSD solvers facilitates the implementation of structural topological optimisation within existing designs. There is no such separation in aerodynamics; the aerodynamic shape is intrinsically linked to the rest of the design by its need to be supported by an underlying structure. This means that aerodynamic topological optimisation of an entire aircraft or wing is unlikely to be a reality in the near or medium term. However there is scope for the aerodynamic topological optimisation of local features of the aerodynamic shape; topological optimisation of wing tips would allow feathered or split winglets of the type seen on the Boeing 737-MAX to be explored much more effectively than at present [22]. An effective topological aerodynamic optimisation framework offers the possibility of radically new designs in applications where aero-structural interaction dominates by leveraging the benefits of both structural and aerodynamic topological changes. In this latter category, applications to Formula 1, strut-braced wing design, engine struts on commercial aircraft and internal engine design could offer significant improvements in performance.

Previous work by the authors of this paper has focused on the development of parameterisation tools for topology optimisation [3, 23] and efficient multi-level schemes for complex aerodynamic optimisation cases [1]. For aerodynamic topology optimisation the r-snake volume of solid (RSVS) parameterisation allows the exploration of arbitrary topology with a single set of volume of solid (VOS) design variables. These let an optimiser transition between the number of bodies inside the flow without having to specify it explicitly. This method has been shown to be effective on topology cases and is capable of exploring shape optimisation cases albeit slightly less efficiently than competing established parameterisation methods. The multi-resolution subdivision scheme is used by Masters et al. [1] in conjunction with an adjoint flow solver and SQP optimiser to produce an efficient framework for multilevel shape optimisation (MLSO). The MLSO framework demonstrated significant improvements in the robustness of the optimisation convergence resulting in the best published results on the first benchmark case of the AIAA Aerodynamic

Design Optimisation Discussion Group (ADODG³). Similar to the multilevel subdivision method is recent work on gradient-limiting shape control [24] whereby surface derivatives are constrained such that smooth and valid iterates are generated. Localised control provides high-fidelity shape-relevant updates, whereas the surface constraints exclude non-physical shapes from the design space; as a result high-fidelity shape optimisation is possible at a reasonable computational cost and future work will incorporate this methodology into the combined framework here.

Previous comparison of the combined optimisation framework has shown marked improvements in all cases compared to its individual components. The goal of this paper is to tackle additional cases with a focus on multi-modal cases using the previously developed and validated framework combining RSVS and MLSO. The combined algorithm has been tested on constrained area supersonic aerodynamic cases [25]; these cases had previously been explored using aerodynamic topology optimisation methods [2, 3]. The study of the aerodynamic behaviour of the current implementation of MLSO has revealed some possible multi-modality even without topological variations. To complement the integration of the parameterisations, optimisers beyond traditional gradient methods and evolutionary algorithms are considered. These include hybrid methods [26], multi-start gradient descents, niching, and illumination methods [27, 28].

II. Approaches to Parametrisation

A parameterisation method capable of tackling aerodynamic topological optimisation using VOS to build aerodynamic shapes has previously been developed by the authors of this study [3]. Similar work within the research group has also led to the development of a very efficient multi-level parameterisation method for aerodynamic shape optimisation. The current work aims to bring the efficiency of the multi-resolution subdivision curves [1] to aerodynamic topology optimisation. This section presents the basics of the RSVS parameterisation and the multi-level subdivision curves and how their integration is implemented.

A. Topological Flexibility using the R-Snake Volume of Solid Parameterisation

Figure 1: R-snake contour (in red) with snaxels (in blue) evolving on the snaking grid (dashed line).

The role of the parameterisation method is to provide an efficient interface between the optimisation method and a solver to form an optimisation framework. This section develops the R-Snake Volume of Solid (RSVS) parameterisation method which blends the topological flexibility of volume of solid design variables with the efficiency of established aerodynamic parameterisation methods. Achieving this level of efficiency requires the RSVS to generate smooth surfaces fulfilling volumes specified on a predefined grid. To ensure the method is flexible enough to support anisotropic design variable refinement and to facilitate the extension to 3-dimensions, the RSVS must be generic enough to work on arbitrary polygonal grids.

The condition used to define the RSVS profile (y) is minimisation of profile length, with the constraint that the area enclosed by the contour within each cell (of contour B_j) must exactly match the specified value for that VOS cell (A_j). This is presented in detail in equation 1. This is analogous to the effect of a tensile force ‘shrink-wrapping’ the required VOS in each cell; the benefit is it allows for smooth profiles in most cases but can also recover sharp corners where the VOS requires it.

³<http://info.aiaa.org/tac/ASG/APATC/AeroDesignOpt-DG>

Figure 2: Volume of solid (VOS) design variables as grey-scale and RSVS profile in red; 1 corresponds to a completely full cell and 0 an empty cell.

$$\begin{aligned} \min \quad & \oint \sqrt{1 + y'^2} dx \\ \text{s.t.} \quad & \oint (y \cap B_j) dx = A_j \quad \forall j \in \{0, \dots, m\} \end{aligned} \quad (1)$$

The R-Snake Volume of Solid (RSVS) parametrisation method relies on a restricted snake (r-snake), a type of parametric active contour, to represent geometric profiles. The r-snake is a closed contour composed of connected vertices called snaxels constrained to travel over the edges of a predefined grid. The movement of these snaxels is governed by a number of simple rules that allow a large range of shapes to be represented and evolved efficiently. The current work uses the parametric active contour developed by Bischoff and Kobbelt in [29] for restricted snakes (r-snakes). An example r-snake with snaxels and snaking grid is shown in figure 1.

The RSVS governing equation (eq. 1) minimises the length of the profile, constraining it to fulfil exactly the volume specified by the VOS design variables. This formulation allows analytical properties of the profiles to be determined using calculus of variations, helping the integration with local and global optimisation frameworks. It can be shown that this minimisation problem leads to the profile being patched arcs of circle, which can themselves be represented by a continuous NURBS. The RSVS profile is built according to the governing equation using a restricted snake (r-snake). The r-snake was chosen as it provides efficient topology handling and is tolerant of any convex layout of VOS design variables. While being independent, the governing equation and the contour recovery method are integrated very efficiently using sequential quadratic programming (SQP).

To define a set of VOS variables a grid is superimposed on the design space, where the design variables become the fraction of each cell within a geometry built from this information. This process is shown for a simple grid in figure 2. This parameterisation procedure provides intuitive handling of topology change without explicit control of it, allowing topological flexibility while maintaining smooth control close to topology changes.

a) Starting RSVS grid and profile (in red) and target geometry (in blue).

b) Final RSVS grid and profile (in red) and target geometry (in blue).

Figure 3: Inverse design of a multi-body aerofoil starting from a 6×2 RSVS grid.

The capability of the RSVS parametrisation has been shown previously on aerodynamic shape and topology cases [3, 23]. A hierarchical approach to RSVS design variables was successfully implemented allowing the exploration of many optimisation cases without any change to the starting position of the optimisation framework [23]. This approach allows the geometric recovery of objects of arbitrary topology, this is shown in figure 3 for a multi-body aerofoil. While this approach is very effective for geometric recovery it is not sufficient to guarantee cheap and robust exploration of aerodynamic design spaces as the refined RSVS grids restrict the relative motions of different geometric features.

B. Efficient Shape Optimisation using Multi-resolution Subdivision Curves

A subdivision scheme defines a curve or surface as the limit of successive refinements starting from some initial polygon or polygonal mesh. Subdivision curves and surfaces currently dominate the entertainment graphics industry due to their unique topological flexibility compared to traditional spline-based methods, however the technology has seen growing attention in engineering applications [1, 30–32]. Recent work by Masters et al. applied multi-resolution subdivision curves in a hierarchical manner to parameterise aerofoil geometry and demonstrated improved efficiency and accuracy of aerodynamic shape optimisation [1, 33]. Whereas the RSVS method provides complete topological flexibility which, in combination with a global search algorithm, also offers excellent coverage of the design space, the multilevel subdivision parameterisation represents an efficient and robust method for precisely resolving the local shape optimum for fixed topology configurations.

Figure 4: Four levels of subdivision of a four point control polygon.

In their work, Masters et al. performed multiple optimisations sequentially, starting from a coarse control mesh and progressively refining; the effect of this is that shape control occurs at different length scales, starting with smooth large-scale changes and progressing to increasingly localised control. In this way high precision shape control can be performed without the deterioration in optimisation efficiency associated with localised shape parameterisation; when used in combination with an adjoint flow solver, providing surface sensitivities at greatly reduced cost, this results in significant reductions in computational cost.

The subdivision formulation is conceptually simple; given an initial control polygon C_0 , a refinement can be made linearly such that a new polygon is derived by a linear relationship using a subdivision matrix P :

$$C_1 = P_0 C_0 \quad (2)$$

This subdivision matrix encompasses two operations: a uniform topological refinement of the mesh (splitting) and a smoothing of the result (averaging), demonstrated in figure 4. Both operations are local and can hence be performed very efficiently. Subdivision schemes with unit maximum eigenvalue converge to a limit surface when applied ad infinitum; in practice the subdivision process can be truncated and the points of the final control polygon can be driven to their final limit positions by a limit matrix P_{eval} . Therefore the limit curve, sampled by the N^{th} subdivision level, can be expressed in terms of the n^{th} level control polygon:

$$\phi_n = P_{eval} P_{N-1} P_{N-1} \dots P_n \quad (3)$$

$$C_\infty = \phi_n C_n \quad (4)$$

Many subdivision schemes exist varying in the properties of the limit surfaces they generate and also, for subdivision surfaces, in the topology of the initial control mesh. Cubic B-Spline equivalent subdivision can be derived using the B-Spline knot-insertion property, this results in the following subdivision matrix for a two-dimensional curve:

$$P = \frac{1}{8} \begin{pmatrix} \ddots & & & & \\ \ddots & & & & \\ & 1 & 6 & 1 & \\ & & 4 & 4 & \\ & & 1 & 6 & 1 \\ & & & 4 & 4 \\ & & & 1 & 6 & 1 & \ddots \\ & & & & & \ddots & \ddots \end{pmatrix} \quad (5)$$

When applied repeatedly, as in equation 3, the columns of the resulting ϕ matrix are cubic B-Spline basis functions.

A powerful extension is the multiresolution formulation, akin to the discrete wavelet transformation, which arises naturally since the subdivision surface definition is factorised as a sequence of refinement operations. By defining a coarsening operator R , in analogy to the inverse of the refinement matrix P of equation 5, then a multiresolution analysis can be performed by decomposing a fine shape representation into a coarse approximation and a detail vector. The former is a coarser control level, which if refined again produces a smooth approximation to the input; the latter are the high frequency details lost during coarsening. Together the two outputs can be used to recover the original fine level input. A multilevel decomposition is performed by recursively applying the coarsening process resulting in a very coarse shape approximation and a set of detail vectors. It is this formulation that is used to preserve optimum geometry when performing progressive refinement optimisations and that allows any arbitrary input geometry to be represented by the subdivision formulation.

III. Integrated Geometry Parameterisation

The focus of the current work is to bring the approaches of the previous section into a single geometry generation framework to allow the efficient exploration of design spaces with a high level of multi-modality. By bringing together the most efficient methods for aerodynamic topology and shape optimisation this new framework significantly reduces the computational cost of existing optimisation cases and improves upon the flexibility and exploration of new types of aerodynamic optimisation problems.

A novel contribution of this work is in the linking of the parameterisation methods and the handling of the two concurrent sets of design variables in a single optimisation process. This is made possible by the multi-resolution subdivision formulation, described in the previous section, which allows any input geometry to be represented in the multiresolution framework (see figure 5). Since the curves generated by RSVS are NURBS, both the RSVS and subdivision representations share a spline equivalence, this is likely to avoid large features in the detail terms of the multi-resolution decomposition allowing efficient transition between the two design spaces. The design geometry can therefore be deformed by both the RSVS design variables (A_j of equation 1) and a coarse set of C_n control points with seamless transfer of information between the two. These two sets of design variables can then be controlled sequentially or simultaneously to achieve topological flexibility with a compact design space (see algorithm 1).

Figure 5: Transfer of geometry between parameterisations.

Algorithm 1 Integrated Parameterisation Optimisation Framework

```
for Each global optimisation step do  
  Generate RSVS Profiles  
  Re-parameterise using multi-level subdivision curves  
  Perform local (topology invariant) optimisation  
  Translate subdivision deformations onto RSVS design space  
  Generate new population of RSVS design variables using a global optimisation process  
end for
```

A. Geometry Interface

Transfer of shape information from the RSVS contours to the subdivision parameterisation is done via high fidelity discretisation of the body profiles. Each RSVS case may contain multiple bodies, each of which is re-sampled to be represented by a piecewise linear ‘loop’. The set of loops provides a general representation of the geometry which allows modular interfacing with the local shape parameterisation method.

Re-sampling the RSVS contours requires smooth surface point distributions with appropriate clustering around sharp corners. Unlike traditional aerofoil optimisation cases there is no a priori knowledge of what the geometry might look like and where clustering might be required. Sharp features requiring clustering of surface points are identified using the exterior angle at each point, smoothed using a moving average over 3% of the points. This measure allows sharp corners (similar to trailing edges) as well as areas of sustained curvature (similar to leading edges) to be identified as extrema of the function. The region of the profile around each identified feature is then sampled using a cosine distribution parametrised by edge length. This process allows symmetric point distributions around features and regions of consistently low curvature have a low surface resolution reducing the CFD mesh density and computational time.

The re-sampled point distribution is then used for re-parameterisation by subdivision, the first step of which is the definition of a high fidelity cubic B-Spline for each loop. This is required for the generation of the evaluation matrix P_{eval} (equation 3) which allows the surface geometry to be practically evaluated without infinite subdivision. This matrix can be derived from the eigen-basis of the subdivision scheme or alternatively from a B-Spline curve defined on the final subdivision level, this latter option is used here. To enable preservation of sharp features, every loop is split into continuous regions between sharp points which are each fit with simple splines by least-squares. The resulting splines are consequently merged on each loop to give the *base* spline. The base spline forms the highest subdivision level and provides the starting point for reverse subdivision. After reverse subdivision, the parameterisation matrices ϕ_n and detail vectors $\mathbf{d}_n$ at each subdivision level are merged across loops to form the multilevel representation required for shape optimisation. This process is illustrated for a three body case in figure 6.

Figure 6: Automated generation of subdivision parameterisation.

B. Exact Well-Posed Shape Control

As presented in section II, a practical feature of the multiresolution subdivision representation is the inclusion of error terms to allow exact recovery of geometry. Common practice for spline-based shape parameterisation is to drop the error term since for an appropriate number of control points the approximation is usually sufficient. However in this work exact transfer of geometry between shape methods is highly desirable to avoid introducing ambiguity into the shape definition. In this work an alternate methodology for exact shape recovery is used whereby the error terms are included into the parameterisation as extra basis functions such that their amplitude can be modified as required by the search algorithm. This has been shown to overcome the adverse effects of including a constant error term into the shape definition [25]. Whereas the error terms are not ‘good’ basis functions in terms of orthogonality and smoothness, this is not of concern when used as part of a multilevel optimisation since the error components are progressively transferred to the standard subdivision basis when refinement is performed. The augmented linear parameterisation is therefore given by:

$$\begin{pmatrix} \mathbf{C}_\infty^x \\ \mathbf{C}_\infty^y \end{pmatrix} = \begin{pmatrix} \phi_n & \mathbf{0} & \mathbf{e}_n^x & \mathbf{0} \\ \mathbf{0} & \phi_n & \mathbf{0} & \mathbf{e}_n^y \end{pmatrix} \begin{pmatrix} \mathbf{C}_n^x \\ \mathbf{C}_n^y \\ \beta^x \\ \beta^y \end{pmatrix} \quad (6)$$

where

$$\begin{pmatrix} \mathbf{e}_n^x & \mathbf{e}_n^y \end{pmatrix} = \sum_{i=n}^N \phi_{i+1} \mathbf{d}_i \quad (7)$$

In addition to the augmented shape parameterisation basis, geometric constraints are also required such that full advantage can be taken of the subdivision parameterisation for efficient local optimisation. This is because subdivision curves, and all spline-based shape representations, are simple shape functions in parametric space and do not include any consideration of the underlying shape connectivity; this is to say that the resulting design space naturally contains non-physical shapes and invalid shapes (e.g. oscillatory, intersecting, inverted etc.), i.e. the problem is under-constrained. Unlike typical academic problems where the initial geometry is known and geometric constraints can be specified manually, the constraint definitions here must be extremely versatile such that they can be applied in an automated fashion to the wide variety of shapes generated by the global topology search. Similarly a balance needs to be struck between sufficient constraint for well-posed search directions and sufficient feasible design space for local exploration. This is achieved in this work through linear constraints implementing move-limits, and a non-linear constraint for surface mesh validity.

IV. Automated Optimisation Framework

In this section, detail is given regarding the methodology for performing aerodynamic optimisation and the individual components that make up the optimisation framework. Section A presents the global and local optimisation methods, followed by a description of the hybrid scheme used to combine global and local search methods. The increased complexity of this hybrid approach arises due to the unsupervised initialisation and running of gradient-based sub-optimisation problems. The implication of this is that reliable evaluation of the flow and sensitivities is essential; the tools and techniques used to achieve this are presented in sections B and C.

A. Optimisation Methods

1. Local Optimisation using Gradient Based Optimiser

The SNOPT [34] (Sparse Non-linear Optimiser) package is used here for gradient-based optimisation. This package implements a Sequential Quadratic Programming (SQP) algorithm for solving general non-linear constrained optimisation problems. The power of this package lies in its ability to efficiently and robustly handle large problems (≈ 1000 s of variables and constraints) while allowing precise constraint satisfaction. The SQP algorithm operates iteratively whereby successive search directions are found from the solution of a quadratic programming (QP) sub-problem and a line-search is used to determine step length. The sub-problems are formed from quadratic approximations to the augmented objective function (Lagrangian) and linearisations of the constraints. The quadratic approximation is initialised with an identity matrix and BFGS updates are used to approach the Hessian of the Lagrangian.

The SQP gradient-based algorithm used in combination with multiresolution subdivision curves for shape parameterisation and an adjoint flow solver (see sec. B) for objective sensitivities, results in an efficient and effective tool for multilevel shape optimisation (MLSO). The MLSO performs sequential shape optimisations starting from a low

fidelity subdivision curve, consisting of few control points, and progressively increasing. After each intermediate optimisation level there is exact transfer of the optimum result from the previous level to the starting geometry for the next level.

2. Global Optimisation by Differential Evolution

Differential Evolution is a heuristic global optimisation method proposed by Storn and Price [35], it was selected due to its robustness and ease of implementation both in serial and parallel. Unlike other heuristic methods it requires few internal parameters and has shown good results on a range of applications [36] including aerodynamic [37] and topology [23] optimisation. This method follows a similar process to Genetic Algorithms (GA) with differences in the crossover and selection stages. DE drives the solution towards the global optimum by combining members of a population; this process follows three stages: combination, crossover and selection. The process used in this work was unchanged from the formulation by Storn and Price [35], the reader is referred to their work for details of the process.

3. Hybrid Optimisers

In most cases of aerodynamic optimisation efficient approaches have either relied on optimisation algorithms whose convergence is linearly correlated with number of design variables [38] or used very low number of design variables with global optimisers [39]. Neither of these approaches is satisfactory for the cases to be tackled by the integrated parameterisation method; the topological flexibility leads to a multi-modal design space which cannot be sufficiently explored using local optimisers and requires too many design variables for the routine use of a global optimiser.

To tackle these shortcomings three procedures were considered: multi-start (MS) local, sequential global to local (SGL), and a hybrid optimisation approach. The multi-start local optimisation uses a randomly generated sample of profiles and optimises all of them, extracting the best result. The sequential approach relies on a set number of DE steps, then uses this population to begin a MS gradient-based optimisation method. The hybrid optimiser is similar to that used by Chernukhin and Zingg [26]. It relies on performing a few local optimisation steps on each member of a global optimisation population before calculating new generations. This approach was shown to be effective in cases where a high number of local minima are present [26]. These three approaches are illustrated in figure 7. Here the gradient-based shape optimisation, performed by MLSO, is identified as the *Local* search, and the topological optimisation, performed using RSVS with DE, is identified as *Global*. For the sequential implementation, the interaction between the two is only weak whereas for the hybrid method they are more strongly coupled.

Initial tests using the hybrid optimiser have been ineffective in the current multi-parameterisation setup. This is due to the poor aerodynamic potential of many of the profiles generated by DE when using the RSVS. For these profiles a simple flow solution is sufficient: any additional computational effort is wasted. For this reason, a small modification

Figure 7: Algorithms for combined global topology and local shape optimisation.

to the hybrid optimiser is proposed: instead of performing MLSO on every DE step, the global optimiser will be allowed to perform more than one DE step before calling on the MLSO framework for accelerated convergence.

B. Flow Analysis and Discretisation

Aerodynamic optimisation relies on the successful integration of CFD, parameterisation and an optimiser into a cohesive framework. In this work a number of options have been implemented to reflect the need for efficiency and topological flexibility. To ensure robust topological flexibility, the RSVS parameterisation method was coupled with a cutcell mesh generator and an unstructured Eulerian flow solver. In order to exploit the topological flexibility of the parameterisation all elements of the optimisation method need to support profiles made of an arbitrary number of bodies. Cutcell mesh generators provide the required flexibility with sufficient accuracy at a low computational cost [40, 41]. The flow solver is an inviscid, compressible unstructured code based on the cell-centred approach by Jameson [42] and following the implementation of Eliasson [43]. The cut-cell mesh generator and flow solver were used in previous studies by Hall et al. [44]. A mesh convergence study was performed on the drag value of the NACA 0012 the converged value of 469.3 drag counts is within 0.3 counts of previous studies using different solvers [33, 45]. The cutcell meshes were used in conjunction with the DE optimiser.

To exploit the efficiency and precision of the multilevel subdivision parameterisation, an adjoint flow solver is required since the high-resolution subdivision levels require surface sensitivities resolved to a correspondingly high-fidelity. For this work, the Stanford University Unstructured (SU²) [46] flow solver is adopted. This open-source software is developed around the task of aerodynamic optimisation and hence has both continuous and discrete adjoint implementations [47]. The main flow solver implements both compressible Euler and RANS equations using an unstructured finite volume method. Multigrid acceleration is available as well as MPI parallel processing. The SU² suite also includes other modules for tasks such as shape parameterisation, mesh adaption and mesh deformation, however only the flow solver module is used here for obtaining flow solutions and flow sensitivities. The continuous adjoint is used here; in SU² the continuous adjoint implementation is formulated on the boundary and hence provides flow sensitivities with respect to infinitesimal perturbations in the local normal directions.

In previous work on topology optimisation for aerodynamics, cut-cell meshes had been used [4, 23, 48]; however these are not compatible with SU², and were replaced with triangular meshes. Topologically flexible mesh generation is performed in an automated manner using Triangle implemented by Shewchuk, a robust, light-weight mesher using constrained Delaunay Triangulation [49]. This method allows bounds to be placed on internal mesh angles and hence provides a cheap and robust method for generating acceptable flow grids across the large variety of shapes and topologies to be encountered. Projected convex hulls of the profile geometry are used to define and control regions of mesh density which decreases with distance away from the surface mimicking the refinement behaviour of cut-cell meshes. Figure 5 shows a close-up of a resulting triangular mesh around a three-body case.

Table 1: SU2 Configuration

Physical problem	Compressible Euler	Continuous adjoint (C_D)
Convective method	Jameson-Schmidt-Turkel	
Time integration	Euler implicit	Runge-Kutta explicit
Artificial Dissipation (k_2, k_4)	0.75, 0.03	2.0, 0.08
Target Residual	10^{-10}	10^{-10}

Table 1 presents the solver settings for SU². These were devised by checking the convergence of the flow solver and the adjoint solver on a hundred profiles generated by the RSVS and meshed by Triangle. These settings maximised the number of converging flow and adjoint solutions. The dissipation on the adjoint flow solution is critical to ensuring reliable convergence on the triangular meshes used in this work.

During shape optimisation, mesh deformation is used to produce new meshes for the displaced surface geometry from the initial volume mesh. Not only is this computationally cheaper than regenerating a mesh for each geometry iteration but it also maintains consistency of the discretisation error which is highly desirable during iterative numerical optimisation. In this work interpolation using multiscale radial basis functions (RBFs) [50] is used. Interpolation using radial basis functions (RBFs) has recently become a prominent mesh deformation method boasting excellent robustness and quality-preserving characteristics [51–53]. The multiscale formulation varies the length scale of the interpolant such that the system solution and update steps are cheap and well-conditioned, while still recovering the exact surface displacements [50].

C. Unsupervised Optimisation

The hybrid optimisation method demands a high degree of reliability since it involves the unsupervised generation and testing of a large number of cases; this inevitably places pressure on the geometry and flow handling methods which need to be exceptionally reliable. For the geometry parameterisation this means fully automated parameterisation and sufficient constraint of the resulting design space, as presented in section III. For the flow analysis, a simple approach is to ensure all errors are caught and handled. However, to maximise the effectiveness of the hybrid optimisation process and achieve meaningful results additional fallback routines can be introduced:

- **Initial objective:** for hybrid optimisation, the initial objective value is recorded and returned in the event of a failure during local shape optimisation such that meaningful information can still be attributed to the profile for further DE iterations;
- **Adjoint contingency:** robustness to failed flow solutions is implicitly provided by the line-search procedure of gradient-based methods; however, the same cannot be said for failed sensitivity calculations. In the event of an adjoint failure (non-convergence) routines are able to adjust numerical parameters such as time-stepping method, CFL number and artificial dissipation in an attempt to obtain a converged adjoint solution;
- **Re-meshing:** the availability of a flexible mesh generation method also allows the possibility of in-loop re-meshing of geometries when large shape deformations have reduced the quality of the grid. Re-meshing is triggered by a pre-set geometric step size or multiple partial/failed convergences of the flow or of the adjoint.

V. Results

The following section presents the results of combined topology and shape optimisations performed by the integrated framework. The integrated optimisation method aims to improve the quick exploration of topological design spaces. Previous aerodynamic topology optimisation frameworks have been hampered by their reliance on agent based optimisation. This has restricted the complexity of the cases that could be explored with a reasonable amount of computational expense. The cases tackled are supersonic drag minimisations under an area constraint for which the benefits of topology optimisation has been demonstrated [2, 3], important features are discussed in section A. One of the main challenges of the MLSO-RSVS is the automation of the optimisation: to successfully explore the design space a large number of local optimisations need to be started from a wide range of starting geometries. The process of tuning the combined method for efficiency and reliability without biasing it to specific cases is presented in section B. The resulting aerodynamic behaviour is studied in detail for a few selected cases (sec. D) showing that the desired aerodynamic features are explored. During this process the MLSO process highlighted potential multi-modality in the local optimisation of two body profiles, because of abrupt changes in flow patterns at the optimal geometries. Finally, the results of the combined framework are shown to be superior to the previous optimisations performed using only the RSVS or the MLSO.

A. Drag minimisation for Fixed Area Profiles at Mach 2

Previous research into aerodynamic topology optimisation has explored inviscid, supersonic, constrained area optimisations as the main test-cases for topology optimisation [23, 48]. In these studies, the use of agent based optimisers was required to explore the topological design space. The goal of the combined framework presented in this paper is to improve performance on these cases both in terms of drag value of the optimum design and computational time. These cases are a useful benchmark for the new algorithm against existing methods and results.

Significant research into these cases was carried out in the 1950s using linearised equations for supersonic flow which yielded analytical optimum solutions. In three dimensions, this effort led to the now famous Sears-Haack profile for minimum wave drag [54, 55]. Similar research by Klunker and Harder [56] used non-linear supersonic pressure coefficient relationships to obtain the profile for minimum pressure drag under thickness and volume constraints. The availability of analytical results for these single body cases provides useful benchmarks for non-linear numerical optimisation frameworks.

Supersonic flows are also an excellent test bed for topology optimisation: there exist multiplane profiles where shock interactions produce bodies with no wave drag [57]. The most well known of these is the Busemann biplane first proposed in the 1930s by Busemann [58]. These cases are of particular interest as the multi-body profiles can be built using the RSVS method: they are known cases for which topological flexibility brings significant drag reduction. An example of the flow around each of these three known linear optima is shown in figure 8.

The mathematical programming representation of topological optimisation case is expressed in equation 8. The behaviour of the optimisation is dependant on the area constraint value c_A . An additional constraint is required for

Figure 8: Three types of linear optimum at Mach 2 with an area (c_A) of 0.08

the multi-body cases to ensure that the optimised profile fits inside the region occupied by the Busemann biplane, the maximum height of the profile (Δy_{max}) cannot be larger than the maximum height of a buseman biplane ($\Delta y_{BUSEMANN}$). This is to avoid the optimiser stacking multiple Busemann bi-planes transverse to the flow, which would lead to good drag performance but is a trivial generalisation of the Busemann bi-plane.

$$\begin{aligned}
 \min \quad & C_D \\
 \text{s.t.} \quad & \sum \mathbf{a} \geq c_A \\
 & \Delta y_{max} \leq \Delta y_{BUSEMANN} \\
 & M = 2
 \end{aligned} \tag{8}$$

B. Configuration and Validation

The purpose of the hybrid optimisation framework is to improve the exploration of the topological design space allowed by the RSVS parameterisation. Inherent in the design of the hybrid method is the potential for a variety of different configurations in executing the optimisation. This includes settings such as parameterisation setup, number of iterations, population management, etc. The key challenge in choosing a configuration is to ensure that the local gradient-based optimisation runs reliably and effectively, regardless of input geometry and without human supervision.

Selection of a robust configuration involved the repeated testing of the combined framework with different settings on the same sample of starting geometries. During testing, focus was given to both the effectiveness of the method, to build confidence that the best solution found was close to the global optimum, as well as efficiency. This validation was performed using MS gradient-based on the starting population of a run of DE for an area constraint value of $c_A = 0.12$.

The following areas were identified for informal hypothesis-testing and the following conclusions were drawn [25]:

1. **Subdivision error basis:** *does the novel error treatment presented here perform better during optimisation than existing methods?*

The new error treatment method was benchmarked against the inclusion of the error as a constant term and discarding the error term. Optimisation results from MS gradient-based runs found that the new implementation consistently out-performed the constant error method, in terms of objective function improvement, without introducing untoward behaviour and while still allowing exact recovery of the input geometry.

2. **Efficient local optimisation:** *what is a acceptable number of local iterations to perform at each multilevel optimisation and in what order should subdivision levels be used?*

Validation runs showed that ten major iterations was sufficient for the majority of the multistart population to reach at least 90% of their capability and for this work this is a good trade-off between computational cost and optimisation effectiveness. This approach is adopted since numerical optimality of the local shape problem is insufficient in triggering a terminating condition for each level and absolute convergence of intermediate levels is not required for the multilevel approach since higher levels offer more effective local shape control. A much larger number of iterations is specified for the final level.

3. **Sub-population:** *can the worst-performing profiles be discarded before starting local shape optimisation?*

Investigations showed conclusively that initial objective value is not a good indicator of final objective after local shape optimisation. Whereas the best initial profile makes approximately a 65% improvement, the best final profile is markedly better improving its original objective value by 87% and the best starting objective by 75%. This behaviour occurs throughout the population with 3 of the top 10 final profiles coming from the worst performing half of the population. Therefore no down-selection of DE population members could be performed before running MLSO.

4. **Sufficient global optimisation:** *what is an acceptable number of DE iterations to perform of global optimisation before starting local optimisation?*

Results for this question are presented in section C as no clear quantitative answer was apparent in the previously published results [25].

C. Selection of the number of DE iterations before execution of the MLSO

The main consideration to make for the sequential “global to local” search (fig 7) is the number of DE iterations required before starting local shape optimisation. DE iterations allow good global exploration and a higher quality population with fewer nonsensical solutions. However, too many iterations leads to a lack of diversity in the population because of convergence of the differential evolution process.

As previously demonstrated, local shape optimisation introduces the capability to significantly improve upon otherwise poor performing agents in the global search. and hence the solution of the global search using RSVS and DE may not correspond to a minima in the high-resolution design space allowed by the combined optimisation framework. This is why starting the local shape optimisation from converged DE results may diminish its effectiveness by repeatedly converging onto the local minima surrounding the solution produced by DE. Figure 9 shows the global convergence of the DE topological optimisation where the results of performing multi-start local shape optimisation have been included at different starting populations. Immediately it is evident how the multi-resolution subdivision parameterisation is able to expand the design space and significantly improve upon the objective. As expected, the local shape optimisation achieves improved results when it is started from a more evolved population; especially in terms of the median result of the locally optimised population. The best result is also improved except for the last case starting after 150 DE iterations; this may be attributable to the aforementioned lack of diversity in converged DE populations.

The current population size of 100 for DE was selected to guarantee robust convergence of the DE algorithm itself. By integrating it with the MLSO method the agent based search method can now be tuned to achieve design space exploration without reliable convergence. This lets the population size be shrunk thereby significantly reducing the computational cost with only a minor penalty on the final objective value.

Figure 9: Convergence of hybrid MLSO-RSVS runs to optimise all agents from five starting points in the global search.

Figure 10: MLSO runs for different settings from 1 to 7 for a single body starting geometry.

D. Validation of the Aerodynamic behaviour

To validate the behaviour of MLSO on the starting geometries generated by the RSVS, a single body geometry and a two body geometry are studied in more detail. For the single body profile, the locally optimum shape is known (see sec. A) and can therefore benchmark the current implementation of the local optimisation framework. Similarly, study of the multi-body case aims to answer the question of modality for geometries similar to Busemann bi-planes. Each of the two cases were tested starting at subdivision levels 1, 3, 5 and 6 progressively refining up to level 7. Each level was allotted either 10 or 100 SNOPT iterations to test the effectiveness of the refinement trigger. Both cases were tested at an area constraint value of 0.12 (eq. 8).

The optimum single body profile, constrained to have an area of 0.12, is expected to resemble the truncated ogive developed by Klunker and Harder [56], shown in figure 8b. This is because the reduction in the angle of the leading edge shock more than makes up for the drag generated by the back pressure applied to the trailing edge. This case presents a significant geometric challenge: the control points and the mesh need to go through an 80 deg turn to capture the flat trailing edge. Figure 10b shows that all the validation runs except one manage to capture this feature. The run which does not capture it is a single level run at the first level of refinement; it simply does not have the resolution to produce a flat surface at the trailing edge.

Figure 10a shows the convergence histories of the MLSO validation runs of the single body. These display the stepped convergence that is expected of multi-level parameterisation: each new level of refinement unlocks a portion of the design space, enabling further improvements. The solid lines in this plot were run using only 10 steps at each level, as analysis of the previous section suggested this would be sufficient. This is confirmed, as the final objective was not compromised by this setting: the runs starting at levels 1, 3 and 5 converged in fewer iterations on a similar design by restricting the number of steps at intermediate levels.

While most optimisation runs capture the blunt trailing edge, they do it with varying levels precision. These discrepancies in geometry are reflected in the final drag values shown in figure 10a. The MLSO appears to struggle to explore this region of the design space and further analysis indicates that this may be a unique instance of the lower subdivision levels producing undesirable starting points for subsequent levels; specifically, attempts by the low fidelity subdivision levels to reproduce the blunt trailing edge have resulted in shapes with concavities which subsequent levels in the multilevel optimisation cannot improve upon. This demonstrates how smooth low dimension shape control is not guaranteed to produce optima in the region of that obtained by higher fidelity control. This issue is further exacerbated by an under-constrained chord length and inaccurate design sensitivities at that location. Future work to incorporate gradient-limiting methods [24] for local shape control is expected to overcome this limitation since these methods are able to provide smooth high-fidelity control without the need for multilevel optimisation. As such, the blunt trailing edge will be contained within the design space initially and spurious low fidelity approximations won't arise. Moreover, the use of gradient-limiting shape control within a single optimisation allows better advantage to be taken of the SQP method, since the Hessian approximation is not lost between successive multilevel optimisations, and as such fewer overall major iterations are required.

Figure 11: MLSO runs at multi-level subdivision settings from 1 to 7 for a starting geometry composed of 2 bodies.

E. Exploration of the Modality of a Two Body Profile

The same testing process was used for a starting geometry made of two bodies. This region of the area constraint space is interesting aerodynamically as the Busemann bi-plane is not optimal. If the shock cancellations seen in figure 8a do not hold, the flow is choked: a large bow-shock forms in front of the geometry causing a step increase in drag for the biplanes at areas above 0.1. To succeed the optimiser will need to balance the choking of the flow with external curved edges which generate shocks that are not cancelled out.

Figure 11 shows that all but the seventh single level case successfully shift the external parts of the geometry to the region between the two bodies. This movement in-board reduces the component of wave drag that is not cancelled by interaction with the second body. This behaviour, shown in figure 11a, appears in the initial steep decrease in drag coefficient down to 0.014.

Despite its desirable initial behaviour, the MLSO does not appear to converge reliably on the same drag value. These differences in drag are reflected in the shape of the throat of the profile, shown in figure 11b. The shape of the profile of the throat is responsible for two aspects of the flow: the quality of the shock cancellations and whether the flow chokes. To understand the reason for this discrepancy, the drag values in the line search direction are plotted,

Figure 12: Analysis of the aerodynamic design space at the optimum geometry, it lies close to the choke point where the flow changes abruptly.

in figure 12a, for the best run in figure 11b. This figure shows a very large discontinuity in the drag value along the SQP search direction; C_D plots on both sides of the discontinuity (fig 12b) reveal the complete change in flow behaviour responsible for this discontinuity. The reason MLSO stalls when it reaches that point in the design space is that information about this change in shock pattern is not captured by the adjoint and cannot be reflected in the search direction. This instability between flow solutions at the optimum solution is similar to the hysteresis with mach number observed in the ADODG case ^{1b} by Destarac et al. [59].

In its current format the final solution to the multi-body shape optimisation problem is very sensitive to starting geometry and parameterisation settings. This is because the need to reduce the external shocks rapidly drives the geometry to the limit of choking. Because choking is caused by the flow going subsonic it is very sensitive to the shock patterns between the bodies. Small differences in paths through the design space lead to this “choking boundary” to be encountered by different profiles. Once this boundary has been encountered the optimiser cannot progress as the existence of the discontinuity is not reflected in the design sensitivities and therefore the search direction.

In its current formulation this case appears to be multi-modal. For this case to be tackled more effectively, additional information needs to be passed to the SQP method informing it about the discontinuous behaviour before it happens. One of the approaches being considered is to add an additional non-linear constraint requiring the minimum local Mach number to be above 1. This would effectively provide the SQP with the direction of the choking boundary once profiles get close to it, allowing the search direction to be adjusted parallel to the boundary so that improvements can still be made.

To explore this case thoroughly changes to the optimisation methods are also being considered. The behaviour of the line-search will be altered to discard steps which are beyond steep discontinuities like the one observed in figure 12a. A change of optimiser is also being explored as an alternative: niching can be used to recover the different local minima existing along this discontinuity, using a small set of the design variables of the throat.

F. Comparison to Previous Aerodynamic Topology Results

The previous sections have presented the data used to validate the convergence behaviour of the combined shape and topology optimisation framework, and highlighted some potential pitfalls. The results in this section shows the performance improvement enabled by the combined MLSO-RSVS framework compared to previous results on the same test cases. The combined framework is compared to the linear theory results discussed in section A and the optimisation results generated by each of the MLSO and RSVS methods in figure 13.

These improvements show the additional capability of the framework compared to the RSVS and MLSO; it opens up the use of the combined method to tackle more complex aerodynamic topology optimisation problems.

^b<http://info.aiaa.org/tac/ASG/APATC/AeroDesignOpt-DG>

Figure 13: Comparison of the combined shape and topology optimisation framework to the linear theory optima and the results of normal ASO and ATO frameworks.

G. Study of optimisation cases with additional topological flexibility

Having shown the capability of the hybrid framework on single and multi-body profiles, additional topological flexibility was added to the design space. This is achieved by increasing the resolution of the RSVS in the vertical direction. For previous cases the RSVS design variables were in a 6 by 10 layout, while these new cases are generated using a 2 by 20. The expectation was that this additional flexibility will allow the multi-start and the sequential optimisation processes to compare the performance of different types of multi-plane profiles more thoroughly. The area constrained drag minimisation case presented in section A is tackled for a constraint value of 0.12.

The first challenge introduced by the additional topological flexibility is the generation of a good starting population. The quality of this population impacts both the multi-start gradient based algorithm and DE. For effective global optimisation a starting population will need sufficient diversity of meaningful aerodynamic profiles. Either of these properties on its own is not sufficient: to adequately explore a fully random population a prohibitive number of agents would be required; and without diversity global optimisation will not be able to generate sufficiently varied designs. The refinement of the design space in the normal direction of the flow led to poor performance in the starting population: many more of the generated profiles are in the choked flow regime. Choking of the flow by far dominates the drag performance of a geometry. In addition, this disproportionately affects profiles made of more bodies. This limitation of the starting population causes the failure of the differential evolution as it rapidly converges on the non-choked profiles that appear most commonly despite their poor drag performance compared to results obtained in previous sections. Significantly increasing the size of the population would allow the starting population generation to generate better profiles and increase the chance of good multi-planes to appear in the population before premature convergence can set in.

Despite the poor aerodynamic performance of the starting population, the multi-start algorithm highlights some interesting behaviours. The severe discontinuity in the aerodynamic design space caused by the change in flow topology poses as challenge: there is no guarantee that drag reduction on the ‘high drag’ side of the discontinuity will reliably guide the optimiser towards the desired change in flow topology. This is especially true for profiles with a more complex topology because of the increased geometric complexity of the local design space at the lowest subdivision level. This can cause the optimiser to follow design directions away from the low drag behaviour allowed by shock cancel-

Figure 14: Pressure coefficient flood plots for some of the profiles optimised by the MS-MLSO framework on a population with increased topological flexibility. In each subfigure the initial drag (C_{D0}) and the final drag (C_D) is stated. The starting profile is inset at the bottom left of the images.

lations. These directions manifest themselves in the very different final geometries generated by the local optimiser from similar starting profiles. Cases are presented in figure 14 for 5 and 6 plane profiles. The best and worst cases presented (figs 14a and 14c respectively) show extremely similar starting profiles with drastically different optimised results. Comparing figures 14a and 14b two profiles with the same starting topology but different geometries converge on profiles relying on very different shock interactions to achieve a low drag value. This suggests multi-modality of the optimisation problem even for a given topology. Interestingly, figure 14d shows a profile with a different topology (6 bodies) but tending to a flow pattern similar to figure 14b. While mostly of academic interest, these cases resemble the shapes of supersonic engine inlets, where shocks are used to perform parts of the compression.

These observations suggest two possible approaches for subsequent studies of this class of optimisation cases: continued search for lower drag; or an attempt to find and classify the local optima that are discovered. A fully hybrid framework, where DE takes into account gradient information to rank its agents, would make up for the shortcoming of the individual optimisers and would lead to better solutions than those found so far for the cases involving more topology. Indeed, by including a measure of the ‘optimisation potential’ of the geometry, the global optimiser, instead of optimising for drag, will optimise the population for its suitability for the MLSO process. The high level of multi-modality both in terms of flow behaviour and geometric topology makes this case ideal for niching and quality diversity [27] approaches. A recent study using a niching variant of differential evolution has successfully been used to identify multiple minima during the optimisation of a wing [60]. These methods would allow convergence and identification of multiple minima.

VI. Combined Shape and Topology Optimisation in Three-Dimensions

In this section, consideration is given to the potential for combined shape and topology optimisation in three-dimensions. This naturally follows the current work since methods for three-dimensional topological representation have been developed within the research group alongside novel approaches to efficient shape optimisation [24]. The concept of topology in engineering design is distinct and more restrictive than that of algebraic geometry due to the importance of features such as sharp edges and geometric constraints; the presence of such discrete features distinguishes geometries of otherwise equivalent mathematical genus. As already discussed, three-dimensional shape control for aerodynamic optimisation is a mature field currently dominated by splines of surfaces (B-Splines, NURBS) and volumes (Bezier, RBF) but is fundamentally restricted in its ability to alter design topology. Specifically, no such method can introduce arbitrary surface features or transition between multiple bodies with a single set of design variables. The previously listed examples of design topology in aerodynamics (wing-tip devices, split wing-tips, strut-braced wings, internal engine design, and Formula 1) all feature performance which is critically linked to the topology of the design and would hence benefit from topological optimisation. As already demonstrated, topological optimisation of external aerodynamics benefits greatly from an incorporated local shape optimisation and this is expected to be all the more significant in three-dimensions where both the topological design space and local shape design spaces are greatly expanded.

A. Three-dimensional topological design

The natural extension of the two-dimensional RSVS parameterisation to three-dimensional surfaces has been formulated: the three-dimensional restricted surface volume of solid (3D-RSVS) parameterisation defines profiles by the minimum surface area that matches the volume fractions specified on a fixed grid. The r-surface method generalises

Figure 15: Coarse wing represented using 60 VOS cells in [2, 5, 6] layout. The colour in the background present the level of convergence of the r-surface on the correct volume.

the connectivity rules of the 2D r-snake such that smooth surfaces for arbitrary topologies can be produced reliably from a compact set of design variables defined on the volume of solid grid. Methodology also exists for the automated design of volume of solid grids which scales to three-dimensions.

B. Robust and flexible shape optimisation

Existing shape control methods can suffer from a manual setup process to define a spline basis and often the parameterisation is constrained such that only one ordinate is free to move or that certain degrees of freedom (*e.g.* sweep, taper) are manually specified. Both of these are not general since they rely on user input and experience which, in the general case, is not available. In addition to geometric restrictions, existing shape control methods often *a priori* constrain the design space at low fidelities and perform poorly when the shape control is refined due to insufficient surface regularisation or large increases in the conditioning of the problem.

Recent work within the research group has developed a new approach to shape optimisation [24] which addresses these issues by performing localised shape control in all coordinate directions and constraining the larger resulting design space by utilising a geometric constraint on surface gradients to maintain validity and continuity of the discrete surface. The gradient-limiting approach can be used with grid-point shape control and therefore does not require the manual configuration of a spline basis. Similarly, the gradient constraints combined with shape control in all coordinate directions mean that the design space is not *a priori* constrained and that shape-relevant perturbations can be made normal to the local shape at every iteration. The invariance of the gradient-limiting formulation to the underlying shape and topology while still being flexible enough to provide useful degrees of freedom means that it is especially applicable to use with topological optimisation in three-dimensions. As in the 2D case, the transfer between the global topological design space and the local shape control space needs to be automatic, since it occurs *in-loop*, and robust, for the large variety of differing geometries produced by the global topological search. The use of the novel gradient-limiting shape control methodology is advantageous in this respect, since it is capable of operating on the grid-points of the analysis geometry and hence does not require the construction of a parametric spline surface or volume.

VII. Conclusions and Future Work

Two optimisation design methods, implementing efficient local shape optimisation and flexible global topology exploration, have been described and demonstrated, and a framework for incorporating the two has been outlined. This framework extends the capability of topological design space exploration by including an efficient local shape representation. The framework consists of: a multi-resolution subdivision shape parameterisation used in combination with an adjoint solver and gradient-based optimiser for robust high-fidelity local shape optimisation; and a restricted snake volume of solid parameterisation and differential evolution search algorithm for flexible coverage of a global topological design space.

Three approaches to integrating the global and local methods have been considered: a multi-start (MS) local, a sequential global to local, and a hybrid optimisation approach, which performs a few local optimisation steps on each member of a global optimisation population before calculating new generations. The main challenge in effectively combining the two methods surrounds the automated and robust transition from the global topology problem to the local shape problem, and solutions to this have been developed. The combined framework has been configured and validated using a test-based methodology to maximise effectiveness and efficiency. As part of this, further insights have been gained on the use of multilevel subdivision curves and the behaviour of population-based gradient optimisations.

The combined framework is benchmarked on an area-constrained supersonic drag minimised problem. This aerodynamic optimisation problem exhibits multi-body solutions and multi-modality and in this work is also shown to pose a significant challenge for gradient-based shape optimisation due to large localised deformations and objective function discontinuities. Multi-modality within a single topology was highlighted adding to the complexity of the design of multi-plane profiles. The aerodynamic behaviour of the multilevel shape optimisation has been validated in detail and is shown to recover expected single-body and two-body profiles. When compared to the two individual topology and shape optimisation methods, the combined framework performs notably better. In addition to expanding the RSVS design space, the inclusion of the MLSO also overcomes the slow convergence of global optimiser necessary to explore topology. These benefits are expected to translate to three dimensions where the problems of design space complexity and exploration are even more acute.

Ongoing work

- **Optimisation under localised constraints:** Work is ongoing to resolve this problem using the hybrid methods demonstrated in this paper; by exploring the constrained design space appropriately, global optimisation should

should require fewer iterations overall and a smaller population size.

- **Niching and quality diversity approaches:** instead of targeting convergence to a single local optima, these methods aim to return a number of locally optimal solutions and forms a natural extension to the combined framework developed here. This allows designers to gain a better understanding of the design space available and can be used at earlier stages of design to suggest design directions. These could be particularly interesting when looking at complex sets of constraints interacting with the topological freedom afforded by the RSVS.

Acknowledgements

This work was carried out using the computational facilities of the Advanced Computing Research Centre, University of Bristol - <http://www.bris.ac.uk/acrc/>. The authors gratefully acknowledge the funding from the EPSRC Research Council UK.

References

- [1] Masters, D. A., Taylor, N. J., Rendall, T. C. S., and Allen, C. B., "Multilevel Subdivision Parameterization Scheme for Aerodynamic Shape Optimization," *AIAA Journal*, aug 2017, pp. 1–16.
- [2] Hall, J., Rendall, T., and Allen, C., "A Volumetric Geometry and Topology Parameterisation for Fluids-based Optimisation," *Computers & Fluids*, Vol. 148, 2017, pp. 137–156.
- [3] Payot, A. D., Rendall, T., and Allen, C. B., "Restricted Snakes: a Flexible Topology Parameterisation Method for Aerodynamic Optimisation," *55th AIAA Aerospace Sciences Meeting, AIAA SciTech Forum, (AIAA 2017-1410)*, American Institute of Aeronautics and Astronautics, Reston, Virginia, jan 2017.
- [4] Payot, A. D., Rendall, T., and Allen, C. B., "Mixing and Refinement of Design Variables for Geometry and Topology Optimization in Aerodynamics," *35th AIAA Applied Aerodynamics Conference, AIAA AVIATION Forum, (AIAA 2017-3577)*, No. June, American Institute of Aeronautics and Astronautics, Reston, Virginia, jun 2017, pp. 1–24.
- [5] Vassberg, J. C. and Jameson, A., "Industrial Applications of Aerodynamic Shape Optimization," *VKI Lectures*, Von Karman Institue, Brussels, 2014, pp. 1–44.
- [6] Sobieczky, H., "Geometry Generator for CFD and Applied Aerodynamics," *New Design Concepts for High Speed Air Transport*, 1997, pp. 137–157.
- [7] Kulfan, B. and Bussioletti, J., "'Fundamental' Parameteric Geometry Representations for Aircraft Component Shapes," *11th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Vol. 1, 2006, pp. 547–591.
- [8] Sobieczky, H., *Recent Development of Aerodynamic Design Methodologies: Inverse Design and Optimization*, chap. Parametric, Vieweg+Teubner Verlag, Wiesbaden, 1999, pp. 71–87.
- [9] Hicks, R. M. and Henne, P. A., "Wing Design by Numerical Optimization," *Journal of Aircraft*, Vol. 15, No. 7, jul 1978, pp. 407–412.
- [10] J. Toal, D. J., Bressloff, N. W., Keane, A. J., and E. Holden, C. M., "Geometric Filtration Using Proper Orthogonal Decomposition for Aerodynamic Design Optimization," *AIAA Journal*, Vol. 48, No. 5, 2010, pp. 916–928.
- [11] Poole, D. J., Allen, C. B., and Rendall, T. C. S., "Metric-Based Mathematical Derivation of Efficient Airfoil Design Variables," *AIAA Journal*, Vol. 53, No. 5, may 2015, pp. 1349–1361.
- [12] Chauhan, D., Chandrashekarappa, P., and Duvigneau, R., "Wing shape optimization using FFD and twist parameterization," *12th Aerospace Society of India CFD Symposium*, Bangalore, 2010.
- [13] Allen, C. B. and Rendall, T. C. S., "CFD-based optimization of hovering rotors using radial basis functions for shape parameterization and mesh deformation," *Optimization and Engineering*, Vol. 14, No. 1, mar 2013, pp. 97–118.
- [14] Vassberg, J. C., Harrison, N. A., Roman, D. L., and Jameson, A., "A Systematic Study on the Impact of Dimensionality for a Two-Dimensional Aerodynamic Optimization Model Problem," *29th AIAA Applied Aerodynamics Conference*, No. 3176, Honolulu, Hawaii, 2011, pp. 1–19.
- [15] Castonguay, P. and Nadarajah, S. K., "Effect of shape parameterization on aerodynamic shape optimization," *45th AIAA Aerospace Sciences Meeting and Exhibit*, , No. January, 2007, pp. 1–20.
- [16] Masters, D. a., Taylor, N. J., Rendall, T. C. S., Allen, C. B., and Poole, D. J., "Geometric Comparison of Aerofoil Shape Parameterization Methods," *AIAA Journal*, Vol. 55, No. 5, may 2017, pp. 1575–1589.
- [17] Masters, D. A., Poole, D. J., Taylor, N. J., Rendall, T., and Allen, C. B., "Impact of Shape Parameterisation on Aerodynamic Optimisation of Benchmark Problem," *54th AIAA Aerospace Sciences Meeting, AIAA SciTech Forum, (AIAA 2016-1544)*, American Institute of Aeronautics and Astronautics, Reston, Virginia, jan 2016, pp. 1–14.
- [18] Hunter, E., "Alternate Detail Part Design and Analysis: Topology, Size, and Shape Optimization of CH-47 Chinook Underfloor Structure," *62nd ANNUAL FORUM PROCEEDINGS-AMERICAN HELICOPTER SOCIETY*, Vol. 1, AMERICAN HELICOPTER SOCIETY, INC, 2006, pp. 260–264.
- [19] Deaton, J. D. and Grandhi, R. V., "A survey of structural and multidisciplinary continuum topology optimization: post 2000," *Struct Multidisc Optim*, Vol. 49, 2014, pp. 1–38.
- [20] Bendsøe, M. P. and Sigmund, O., "Material interpolation schemes in topology optimization," *Archive of Applied Mechanics*, Vol. 69, 1999.
- [21] Rozvany, G. I. N., "A critical review of established methods of structural topology optimization," *Structural and Multidisciplinary Optimization*, Vol. 37, 2009, pp. 217–237.

- [22] Gagnon, H. and Zingg, D. W., "Two-Level Free-Form and Axial Deformation for Exploratory Aerodynamic Shape Optimization," *AIAA Journal*, Vol. 53, No. 7, jul 2015, pp. 2015–2026.
- [23] Payot, A. D. J., Rendall, T. C. S., and Allen, C. B., "Restricted Snakes: a Flexible Topology Parameterisation Method for Aerodynamic Optimisation," *55th AIAA Aerospace Sciences Meeting*, 2017.
- [24] Kedward, L., Allen, C. B., and Rendall, T., "Gradient-Limiting Shape Control for Efficient Aerodynamic Optimisation," *2018 Applied Aerodynamics Conference*, 2018, pp. 1–22.
- [25] Kedward, L. J., Payot, A. D. J., Rendall, T. C. S., and Allen, C. B., "Efficient Multi-Resolution Approaches for Exploration of External Aerodynamic Shape and Topology," *36th AIAA Applied Aerodynamics Conference*, American Institute of Aeronautics and Astronautics, June 2018.
- [26] Chernukhin, O. and Zingg, D. W., "Multimodality and Global Optimization in Aerodynamic Design," *AIAA Journal*, Vol. 51, No. 6, jun 2013, pp. 1342–1354.
- [27] Pugh, J. K., Soros, L. B., and Stanley, K. O., "Quality Diversity: A New Frontier for Evolutionary Computation," *Frontiers in Robotics and AI*, Vol. 3, No. July, 2016, pp. 1–17.
- [28] Gaier, A., Asteroth, A., and Mouret, J.-b., "Aerodynamic Design Exploration through Surrogate-Assisted Illumination," *18th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, AIAA AVIATION Forum, (AIAA 2017-3330)*, American Institute of Aeronautics and Astronautics, Reston, Virginia, jun 2017.
- [29] Kobbelt, L. P. and Bischoff, S., "Parameterization-free active contour models with topology control," *The Visual Computer*, Vol. 20, No. 4, jun 2004, pp. 217–228.
- [30] Shepherd, P. and Richens, P., "The case for Subdivision Surfaces in building design," *Journal of the International Association for Shell and Spatial Structures*, Vol. 53, No. 174, 2012, pp. 237–245.
- [31] Bornemann, P. B. and Cirak, F., "A subdivision-based implementation of the hierarchical b-spline finite element method," *Computer Methods in Applied Mechanics and Engineering*, Vol. 253, 2013, pp. 584–598.
- [32] Bandara, K., Ruberg, T., and Cirak, F., "Shape optimisation with multiresolution subdivision surfaces and immersed finite elements," *Computer Methods in Applied Mechanics and Engineering*, Vol. 300, 2016, pp. 510–539.
- [33] Masters, D. A., Taylor, N. J., Rendall, T., and Allen, C. B., "A Locally Adaptive Subdivision Parameterisation Scheme for Aerodynamic Shape Optimisation," *34th AIAA Applied Aerodynamics Conference*, American Institute of Aeronautics and Astronautics, Reston, Virginia, jan 2016, p. 3866.
- [34] Gill, P. E., Murray, W., and Saunders, M. A., "SNOPT: An SQP Algorithm for Large-Scale Constrained Optimization," *SIAM Review*, Vol. 47, No. 1, 2005, pp. 99–131.
- [35] Storn, R. and Price, K., "Differential evolution: a simple and efficient heuristic for global optimization over continuous spaces," *Journal of global optimization*, 1997, pp. 341–359.
- [36] Mezura-Montes, E. and Coello Coello, C. a., "Constraint-handling in nature-inspired numerical optimization: Past, present and future," *Swarm and Evolutionary Computation*, Vol. 1, No. 4, dec 2011, pp. 173–194.
- [37] Bevan, R., Poole, D., Allen, C., and Rendall, T., "Adaptive Surrogate-Based Optimization of Vortex Generators for a Tiltrotor Geometry," *Journal of Aircraft*, Vol. 54, No. 3, 5 2017, pp. 1011–1024.
- [38] Lyu, Z., Xu, Z., and Martins, J. R. R. A., "Benchmarking Optimization Algorithms for Wing Aerodynamic Design Optimization," *8th International Conference on Computational Fluid Dynamics (ICCFD8)*, Chengdu, China, jul 2014.
- [39] Poole, D., Allen, C., and Rendall, T., "High-fidelity aerodynamic shape optimization using efficient orthogonal modal design variables with a constrained global optimizer," *Computers & Fluids*, Vol. 143, 2017, pp. 1–15.
- [40] Aftosmis, M. J., Berger, M. J., and Melton, J. E., "Robust and efficient Cartesian mesh generation for component-based geometry," *AIAA Journal*, Vol. 36, No. October, 1998, pp. 952–960.
- [41] Rendall, T. and Allen, C., "CFD Simulation of Arbitrary Motion in Two-Dimensional Spacetime Using Cut-Cell Meshes," *28th AIAA Applied Aerodynamics Conference*, No. AIAA 2010-4696, American Institute of Aeronautics and Astronautics, Reston, Virginia, jun 2010, pp. 1–11.
- [42] Jameson, A., Schmidt, W., and Turkel, E., "Numerical solutions of the Euler equations by finite volume methods using Runge-Kutta time-stepping schemes," *AIAA paper*, Vol. n/a, No. n/a, 1981, pp. n/a.
- [43] Eliasson, P., *EDGE: A Navier-Stokes solver for unstructured grids*, 2001.
- [44] Hall, J., Rendall, T., Allen, C., and Peel, H., "A multi-physics computational model of fuel sloshing effects on aeroelastic behaviour," *Journal of Fluids and Structures*, Vol. 56, 2015, pp. 11–32.
- [45] Poole, D. J., Allen, C. B., and Rendall, T., "Control Point-Based Aerodynamic Shape Optimization Applied to AIAA ADODG Test Cases," *53rd AIAA Aerospace Sciences Meeting, AIAA SciTech Forum, (AIAA 2015-1947)*, No. January, American Institute of Aeronautics and Astronautics, Reston, Virginia, jan 2015, pp. 1–20.
- [46] Economou, T. D., Alonso, J. J., Albring, T. A., and Gauger, N. R., "Adjoint Formulation Investigations of Benchmark Aerodynamic Design Cases in SU2," *35th AIAA Applied Aerodynamics Conference*, , No. June, 2017, pp. 1–13.

- [47] Palacios, F., Colonno, M. R., Aranake, A. C., Campos, A., Copeland, S. R., Economon, T. D., Lonkar, A. K., Lukaczyk, T. W., Taylor, T. W. R., and Alonso, J. J., "Stanford University Unstructured (SU2): An open-source integrated computational environment for multi-physics simulation and design," *51st AIAA Aerospace Sciences Meeting*, 2013.
- [48] Hall, J., Poole, D. J., Rendall, T., and Allen, C. B., "Volumetric Shape Parameterisation for Combined Aerodynamic Geometry and Topology Optimisation," *16th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, AIAA AVIATION Forum, (AIAA 2015-3354)*, No. June, American Institute of Aeronautics and Astronautics, Reston, Virginia, jun 2015, pp. 1–29.
- [49] Shewchuk, J. R., "Delaunay refinement algorithms for triangular mesh generation." *Computational Geometry*, Vol. 22, 2002, pp. 21–74.
- [50] Kedward, L., Allen, C. B., and Rendall, T. C., "Efficient and exact mesh deformation using multiscale RBF interpolation," *Journal of Computational Physics*, Vol. 345, 2017, pp. 732–751.
- [51] de Boer, A., van der Schoot, M. S., and Bijl, H., "Mesh deformation based on radial basis function interpolation," *Computers and Structures*, 2007.
- [52] Rendall, T. C. S. and Allen, C. B., "Unified fluid-structure interpolation and mesh motion using radial basis functions," *International Journal for Numerical Methods in Engineering*, Vol. 74, No. 10, 2008.
- [53] Rendall, T. C. S. and Allen, C. B., "Reduced Surface Point Selection Options for Efficient Mesh Deformation using Radial Basis Functions," *Journal of Computational Physics*, 2010.
- [54] Haack, W., *Projectile shapes for smallest wave drag*, Graduate Division of Applied Mathematics, Brown University, 1948.
- [55] Sears, W. R., "On projectiles of minimum wave drag," *Quart. Appl. Math*, Vol. 4, No. 4, 1947, pp. 361–366.
- [56] Klunker, E. B. and Harder, K., "Comparison of Supersonic Minimum-Drag Airfoils Determined By Linear and Nonlinear Theory," Tech. Rep. February, NATIONAL AERONAUTICS AND SPACE ADMINISTRATION HAMPTON VA LANGLEY RESEARCH CENTER, Washington, 1952.
- [57] Roshko A; Liepmann, H. W., *Elements of gas dynamics*, Wiley, New York, 1957.
- [58] Busemann, A., "Aerodynamic lift at supersonic speeds," *Luftfahrtforschung*, 1935.
- [59] Destarac, D., Carrier, G., Anderson, G. R., Nadarajah, S., Poole, D. J., Vassberg, J. C., and Zingg, D. W., "Example of a Pitfall in Aerodynamic Shape Optimization," *AIAA Journal*, 2018.
- [60] Poole, D. J., Allen, C. B., and Rendall, T., "Identifying Multiple Optima in Aerodynamic Design Spaces," *2018 Multidisciplinary Analysis and Optimization Conference*, American Institute of Aeronautics and Astronautics, Reston, Virginia, jun 2018.